AFRL-SR-AR-TR-06-0045

REPORT DOCUMENTATION PAGE

1. REPORT DATE (DD-MM-YYYY) 07-02-2006	2. REPORT TYPE Final Progress Report	3. DATES COVERED (From - To) Jan 1 2003-November 30 2005

4. TITLE AND SUBTITLE
Renewal of the Attentive Sensing Project

5a. CONTRACT NUMBER

5b. GRANT NUMBER
F49620-03-1-0117

5c. PROGRAM ELEMENT NUMBER

6. AUTHOR(S)
Darryl Morrell

5d. PROJECT NUMBER

5e. TASK NUMBER

5f. WORK UNIT NUMBER

7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES)
Department of Engineering

Center Building Room 110
Arizona State University
Mesa AZ 85212-0180

8. PERFORMING ORGANIZATION REPORT NUMBER
EGR-DM-0301

9. SPONSORING / MONITORING AGENCY NAME(S) AND ADDRESS(ES)
Dr. Jon Sjogren/NM
Air Force Office of Scientific Research
875 North Randolph Street
Suite 325 Room 3112
Arlington, Virginia 22203

10. SPONSOR/MONITOR'S ACRONYM(S)
AFOSR

11. SPONSOR/MONITOR'S REPORT NUMBER(S)

12. DISTRIBUTION / AVAILABILITY STATEMENT
Distribution Unlimited

DISTRIBUTION STATEMENT A
Approved for Public Release
Distribution Unlimited

13. SUPPLEMENTARY NOTES

14. ABSTRACT
This report is the final progress report for a project whose goal is to examine a class of problems related to the effective allocation of sensing resources in situations involving tunable sensors, configurable sensor suites, or constraints in communication bandwidth and processing resources. We characterized the performance of optimal and suboptimal sensor allocation strategies for target location and detection problems. We have also implemented and characterized the performance of target-tracking algorithms that use a configurable foveal sensor and a pan-tilt-zoom camera.

15. SUBJECT TERMS
Sensor Configuration, Optimal Bayesian Control, Heuristic Control

16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT	18. NUMBER OF PAGES	19a. NAME OF RESPONSIBLE PERSON Darryl Morrell
a. REPORT UNCLASSIFIED	b. ABSTRACT UNCLASSIFIED	c. THIS PAGE UNCLASSIFIED	UL	13	19b. TELEPHONE NUMBER (include area code) 480-727-1123

Standard Form 298 (Rev. 8-98)
Prescribed by ANSI Std. Z39.18

Renewal of the Attentive Sensing Project
Final Report

January 1, 2003 through November 30, 2005

Darryl Morrell-PI

Department of Engineering
Arizona State University at the Polytechnic Campus
Mesa AZ 85212-0180

AFOSR Grant F49620-03-1-0117

1 Executive Summary

The overall objective of the Renewal of the Attentive Sensing project was to investigate effective allocation and/or configuration of sensing resources in target detection, localization, classification, and tracking problems. In this executive summary, we briefly summarize the technical accomplishments in this project from January 1, 2003 to November 30, 2005. We also describe the personnel supported by this project, the publications resulting from the project, and technical interactions arising from this project.

1.1 Technical Accomplishments

The technical work done in this project was organized into two broad areas. The first was attentive detection and localization of a target using a configurable sensor (Task 1 In the proposal); this area dealt primarily with problems in which the unknowns and sensor configuration variables were discrete. The primary thrust in Task 1 was extending previous work to more complex detection of localization problems and developing theoretical bounds on the performance of detectors. The second area was attentive estimation of the state of a dynamic system (Task 2 In the proposal); this area involved unknowns and sensor configuration variables that were continuous. We investigated applications that involve foveal sensors and pan-tilt-zoom cameras. We also investigated the development of configurable sensors for simultaneous detection and tracking of targets. We briefly describe the progress in each of these two areas.

1.1.1 Detection and Localization of Targets

This work consisted of extending previous work on attentive detection using simple sensors to more complex detection and localization problems. A significant effort was also devoted to the development of computationally tractable bounds on the performance of optimal sensor configuration rules.

We have developed a heuristic configuration rule for a sensor that can be configured to interrogate a block of adjacent cells. This heuristic is motivated by the divide-and-conquer approach that is successful in constrained twenty questions problems. The heuristic uses greedy clustering to divide the grid into two approximately equiprobable regions, and interrogates the region with the most probable cell; the clustering begins with the most probable cell. If the same region is selected in two consecutive time epochs, the sensor interrogates only the most probable cell in the second time epoch. The sensor configuration heuristic was evaluated using Monte Carlo simulations. The performance of the multi-cell sensor is compared to the best-performing single-cell interrogation rule found to date. As expected, the multi-cell sensor out performs the single-cell sensor for both criteria.

We also investigated the development of performance bounds for a single mode sensor. Actually finding bounds requires significant computation, so the primary development effort was finding computational approaches that give good bounds with acceptable complexity. The computation of these bounds usually includes (either directly or indirectly) computation of sensor configuration rules; as the bound becomes tighter, the sensor configuration rule typically approaches the optimal rule. It should be noted that any heuristic decision rule results in either a probability of error or an expected number of sensor uses larger than the optimal; thus, evaluating the performance of heuristic decision rules provides upper bounds on achievable performance. We investigated three basic types of bounds:

1. Application of variational methods to graphical models.
2. Approximation of cost-to-go (value) functions.
3. Decision tree traversal and pruning (*e.g.* branch and bound) methods.

To date, none of these approaches has resulted in satisfactory bounds.

1.1.2 Attentive Estimation of the State of a Stochastic System

We investigated attentive estimation of dynamical systems in the context of target tracking; this work used a foveal sensor as well as pan-tilt-zoom cameras.

We used stochastic optimization to implement open-loop control strategies for the foveal sensor tracking a single target. Since stochastic optimization is too computationally demanding for most applications, we also developed a good, computationally feasible, myopic (greedy) heuristic using the optimal myopic strategy as a model. We applied the optimal and heuristic control strategies in open-loop, open-loop feedback, and approximate closed-loop configuration algorithms. We characterized the resulting tracker performance using Monte Carlo simulation; myopic control rules (both stochastic optimization and the heuristic) gave better tracker performance than multi-step open-loop or open-loop feedback controls. The best performance was obtained using myopic optimization or an approximate multi-step closed-loop control.

In many scenarios, targets may be difficult to detect with a single use of a sensor due to target stealth, clutter, or other factors. The use of a tracker that incorporates a target dynamics model and multiple uses of the sensor may allow better target detection. Such as scheme, in which the estimator jointly tracks and make decisions about target presence or absence, is denoted track before detect. We have investigated joint tracking and detection in the context of the foveal sensor. We adapted the near optimal sensor configuration heuristic to the pixelated foveal sensor model implemented with a particle filter that estimates both number of targets and the target states, and found that it did a good job of tracking targets that had a reasonably high signal-to-noise ratio. We were less successful simultaneously tracking and detecting dim targets; we believe that further development of the particle filter implementation may improve the performance.

We also investigated heuristic strategies to configure a foveal sensor to track multiple targets moving in one dimension. These strategies were obtained by extending the heuristic strategy to multiple targets using three approaches: simultaneously observe all targets, center the foveal region on each target in turn, and center the foveal region on the target with the worst position estimate . Our simulation results showed the best performance is obtained by the last rule.

We adapted the concepts develop for the foveal sensor configuration algorithms to the problem of configuring pan-tilt-zoom cameras to track a target maneuvering in three dimensions. In particular, we developed an adaptive zoom algorithm that minimizes localization errors by adaptively changing the focal length of the camera.

1.2 Personnel Supported

Personnel supported by this project include:

1. Darryl Morrell, PI
2. Fengjun Xi, Research Assistant
3. Himanshu Shah, Research Assistant

1.3 Publications

The following conference papers were written as a consequence of this work:

1. H. Shaw and D. Morrell, "A New Adaptive Zoom Algorithm for Tracking Targets Using Pan-Tilt-Zoom Cameras," *Conference Record of the 36th Asilomar Conference on Signals, Systems, and Computers*, October 2005.
2. F. Xi and D. Morrell, "Tracking Multiple Closely Spaced Targets using an Adaptive Foveal Sensor," (invited paper), *Proceedings of ICASSP'05*, March 2005, pp. v/941–v/944.
3. F. Xi and D. Morrell, "Target Tracking Using an Image Sensor with a Configurable Active Area," *Conference Record of the 36th Asilomar Conference on Signals, Systems, and Computers*, November 2004, pp. 2111–2115.
4. H. Shah and D. Morrell, "An Adaptive Zoom Algorithm For Tracking Targets Using Pan-Tilt-Zoom Cameras," *Proceedings of ICASSP'04*, May 2004, Volume II, pp. II-721–II-724.
5. Y. Xue and D. Morrell, "Target Tracking and Data Fusion using Multiple Adaptive Foveal Sensors," *6th International Conference on Information Fusion*, July 2003, pp. 326–333.
6. D. Morrell and W. Stirling, "An Extended Set-valued Kalman Filter," *3rd International Symposium on Imprecise Probabilities and Their Applications*, July 2003, pp. 396–407.

The following theses were supported by this grant:

1. Fengjun Xi, *Sensor Configuration for Target Tracking* , M. S., April 2005.
2. Himanshu Shah, *An Adaptive Zoom Algorithm for Tracking Targets Using Pan-Tilt-Zoom Cameras*, M. S., December 2003.

1.4 Technical Interactions

The following research interactions were related to this project:

1. D. Morrell presented a talk on particle filters and the foveal sensor on February 21, 2003, to the Signal Processing Lunch and Learn Seminar conducted at General Dynamics.

2. D. Morrell participated as an invited speaker in a panel discussion on sensor networks at the 22nd IEEE International Performance, Computing, and Communications Conference on April 10, 2003.

3. D. Morrell was asked to organize and chair a special session at at the Asilomar Conference on Signals, Systems, and Computers, November 2004.

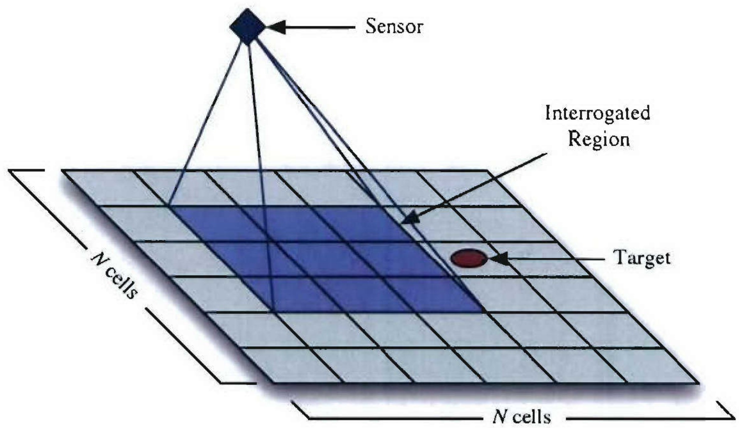

Figure 1: Target localization scenario. The space to be searched is divided into a rectangular grid. At each time epoch, the sensor chooses a rectangular subset of the grid to interrogate.

2 Introduction

The technical work done in this project was organized into two broad areas. The first was attentive detection and localization of a target using a configurable sensor (Task 1 In the proposal); this area dealt primarily with problems in which the unknowns and sensor configuration variables were discrete. The primary thrust in Task 1 was extending previous work to more complex detection of localization problems and developing theoretical bounds on the performance of detectors. The second area was attentive estimation of the state of a dynamic system (Task 2 In the proposal); this area involved unknowns and sensor configuration variables that were continuous. We investigated applications that involve foveal sensors and pan-tilt-zoom cameras. We also investigated the development of configurable sensors for simultaneous detection and tracking of targets. We briefly describe the progress in each of these two areas.

3 Task 1: Attentive Detection and Localization

Task 1 consisted of extending previous work on attentive detection using simple sensors to more complex detection and localization problems. A significant effort was also devoted to the development of computationally tractable bounds on the performance of optimal sensor configuration rules.

3.1 Detection and Localization of Targets

We considered the following abstract two-dimensional target localization problem. An area of interest is divided into a grid containing $N \times N$ cells as shown in Figure 1. A stationary target is known to be located in one of the cells; the prior probability of the target being located in each cell is known.

At each time epoch, the sensor is configured, an observation is obtained, and this observation is used to update the probability distribution of the target location using Bayes theorem. The sensor can select a rectangular group of adjacent cells to interrogate. The result of the interrogation is a binary valued observation. If the target is in the interrogated area, the target will be detected with a known probability of detection; if the target is not in the interrogated area, a false alarm will be generated with a know probability.

Sensor configuration rules were investigated with respect to two optimality criteria:

1. Minimize the probability of error when the sensor is applied a fixed number of times (denoted *MinPE*).

2. Minimize the number of sensor uses needed to achieve a given probability of error (denoted *MinUse*).

For this sensor under the *MinPE* optimality criterion, the myopic optimal rule is to interrogate either the most probable or second most probable cell. This rule does not take advantage of the sensor's multi-cell interrogation capability.

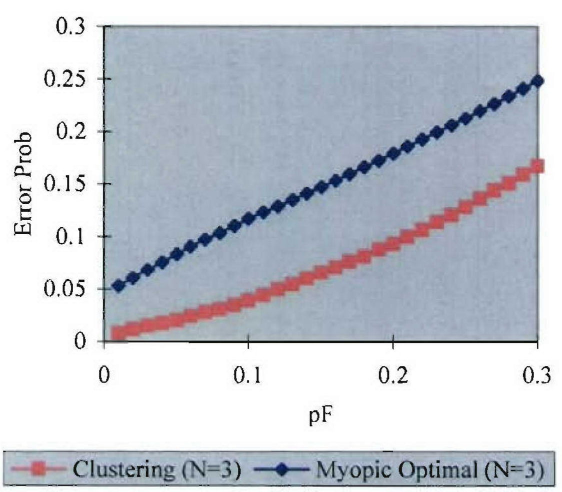

(a)

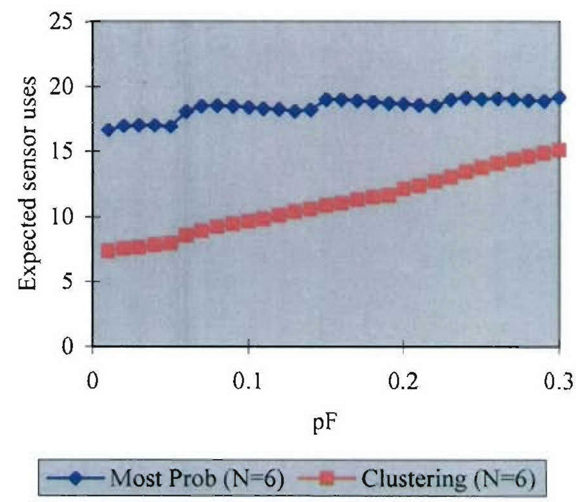

(b)

Figure 2: Performance of the sensor configuration heuristic: (a) comparison of probability of error after 10 sensor uses on a 3×3 grid for the myopic optimal rule and the heuristic. (b) comparison of expected number of sensor uses needed to obtain a probability of correct decision of 0.95 on a 6×6 grid for the rule that interrogates the most probable cell and the heuristic.

We have developed a heuristic configuration rule with better performance (over time) than the myopic optimal rule. This heuristic is motivated by the divide-and-conquer approach that is successful in constrained twenty questions problems. The heuristic uses greedy clustering to divide the grid into two approximately equiprobable regions, and interrogates the region with the most probable cell; the clustering begins with the most probable cell. If the same region is selected in two consecutive time epochs, the sensor interrogates only the most probable cell in the second time epoch.

Computational evaluation of the sensor configuration heuristic for both optimality criteria is shown in Figure 2. The performance of the multi-cell sensor is compared to the best-performing single-cell interrogation rule found to date: the myopic optimal (for the *MinPE* criterion) or choosing the most probable cell (for the *MinUse* criterion). As expected, the multi-cell sensor out performs the single-cell sensor for both criteria.

3.2 Configuration Performance Bounds

We investigated the development of performance bounds for the single mode sensor for both the *MinPE* and *MinUse* optimality criteria. Actually finding bounds requires significant computation, so the primary development effort is finding computational approaches that give good bounds with acceptable complexity. The computation of these bounds usually includes (either directly or indirectly) computation of sensor configuration rules; as the bound becomes tighter, the sensor configuration rule typically approaches the optimal rule. It should be noted that any heuristic decision rule results in either a probability of error or an expected number of sensor uses larger than the optimal; thus, evaluating the performance of heuristic decision rules provides upper bounds on achievable performance.

We have investigated three basic types of bounds:

1. Application of variational methods to graphical models.
2. Approximation of cost-to-go (value) functions.
3. Decision tree traversal and pruning (*e.g.* branch and bound) methods.

In the following, we briefly describe our work on each of these three approaches and outline their potential benefits and drawbacks. To date, none of these approaches has resulted in satisfactory bounds for either optimality criterion.

In variational methods, complex graphical models are approximated with simpler models; variational parameters are introduced to account for model approximations. The nature of the variational parameters determines whether upper or lower bounds will be computed. We investigated these methods briefly, but did not pursue them seriously; our sensor configuration problems are sequential decision problems in which the current decision depends on all previous decisions and observations, and these problems have very highly connected graphical models which do not appear to lend themselves to variational analysis.

The approximation of cost-to-go functions (also known as value functions) has received significant attention from researchers working on Partially Observed Markov Decision Processes (POMDP). In our case, the cost-to-go function is the expected probability of error after completing the remaining observations expressed as a function of the current posterior probability distribution, and applies only for the *MinPE* criterion. In the POMDP literature, several different approximation approaches have been developed, typically with the goal of computing good control rules; the bounds are obtained as a byproduct of this computation. Value functions for POMDP problems are piecewise linear and convex; the complexity of their representation increases exponentially with the number of time steps ahead that the sensor is scheduled. Interpolation of a value function provides a lower bound to the probability of error; elimination of components of the piecewise linear representation can give upper bounds, but choosing which components to eliminate is a computationally difficult problem. We investigated several approximation approaches, including different interpolation schemes and selection of interpolation points; in our work, we found that easily computable interpolation functions give poor bounds.

Traversing the decision tree that represents possible sequences of sensor configurations can provide optimal configuration rules; partial traversals, along with bounds associated with the un-traversed portions of the tree, can provide both upper and lower bounds. The decision tree method is capable of computing bounds for both criteria. Practical implementation of the tree traversal method requires good heuristics to guide the traversal. We have investigated the AO* algorithm, which requires a lower bound estimate of the performance along a given (unsearched) branch of the decision tree.

For problems in which the target characteristics do not change (e.g. localization and detection of a fixed target, classification of a target), the posterior probability distribution is independent of the order in which observations are obtained. Thus, the decision tree has many nodes that correspond to the same posterior probability; upper and lower bounds for only one of these nodes must be evaluated to compute the overall upper and lower bounds. Thus, the decision tree is equivalent to a graph structure with far fewer nodes, leading to a significant computational savings when evaluating bounds. To date, our lower bounds on unsearched branches of the tree have not led to computationally tractable bounds.

4 Task 2: Sensor Configuration for Tracking and Track Before Detect

Task 2 consisted of investigation of attentive estimation of dynamical systems in the context of target tracking; this work used a foveal sensor as well as pan-tilt-zoom cameras. A foveal sensor has a central high-acuity (foveal) region surrounded by a lower-acuity region; the foveal region can be steered and its width can be adjusted. We developed an optimal myopic (*i.e.* greedy) sensor configuration process using stochastic approximation but found that it was unsuited for real time implementation. However, a heuristic configuration rule derived from the characteristics of the optimal ruled provided good performance. We extended this heuristic strategy to multiple targets and to track before detect sensors.

To apply principles developed in the context of foveal sensors to more immediately practical sensors, we adapted the heuristic configuration rules to pan-tilt-zoom cameras. Configuring pan and tilt corresponds to configuring the foveal center; configuring zoom corresponds to setting the gain of the foveal region. Our investigation indicated that the configuration heuristics improve the target tracking accuracy obtainable with a pan-tilt-zoom camera.

4.1 Optimal Rules for Foveal Sensors

We used stochastic optimization to implement open-loop control strategies for the foveal sensor tracking a single target. Since stochastic optimization is too computationally demanding for most applications, we also developed a good, computationally feasible, myopic (greedy) heuristic using the optimal myopic strategy as a model. We applied the optimal and heuristic control strategies in open-loop, open-loop feedback, and approximate closed-loop configuration algorithms. We characterized the resulting tracker performance using Monte Carlo simulation; myopic control rules (both stochastic optimization and the heuristic) gave better tracker performance than multi-step open-loop or open-loop feedback controls. The best performance was obtained using myopic optimization or an approximate multi-step

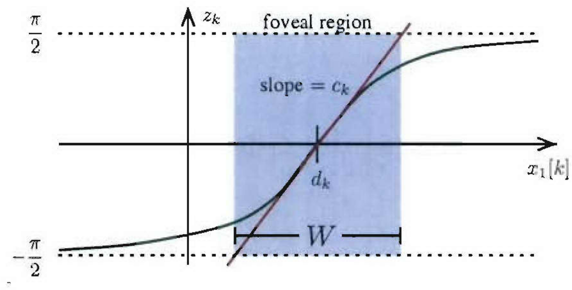

Figure 3: Output function of the foveal sensor.

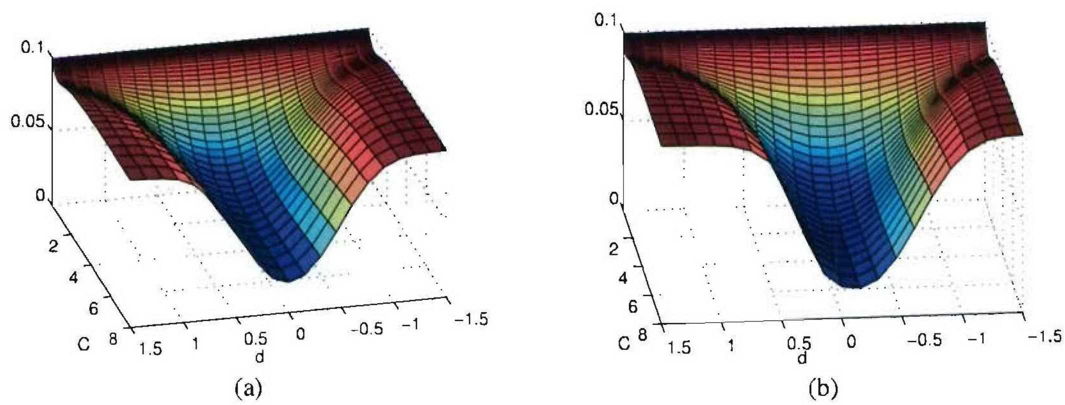

Figure 4: Plots of the expected squared position error for $M = 1$ using (a) a symmetric density for the predicted position, and (b) an asymmetric density.

closed-loop control.

We considered a target moving in one dimension whose movement was modeled using a discrete-time linear system. The foveal sensor provides observations of the target position X_k corrupted by noise. As shown in Figure 3, the foveal sensor has a high gain region surrounded by low gain regions. The sensor has two adjustable parameters: d_k sets the location of the foveal region, while c_k sets the gain at the center of the foveal region. The observation of the target is related to its position as

$$z_k = \arctan\left(c_k[X_k - d_k]\right) + v_k,$$

where v_k is white Gaussian noise with variance R.

We investigated configuration of the sensor from a given time k to a time M steps in the future to minimize the expected squared position error at time $k + M$. Configuring the sensor involves choosing values for the sequence of foveal centers and foveal gains. To provide a clear understanding of the optimization problem, we computed the expected squared position error numerically for $M = 1$ using two different densities for the predicted target position. Figure 4(a) shows this error when the density is Gaussian with mean zero and variance 0.1. The expected squared error function is almost flat for a wide range of c_{k+1} values, indicating that a wide range of values will give near optimal results and that gradient descent optimization will converge slowly. Figure 4(b) shows the objective function when the density is an asymmetric two-component Gaussian mixture with mean zero and variance 0.1. Note that the asymmetry shifts the optimizing value of d_{k+1} to approximately -0.1.

To compute optimizing values for the foveal sensor parameters, we used a stochastic gradient descent approach as follows. In the tracker implementation, the density of the predicted target position is represented using particles and weights computed by the particle filter; no closed form representation of the density is available. Using these particles and weights to generate Monte Carlo samples of future states and observations, we can compute expected squared error values for various settings of the sensor configuration variables. Since this computation is based on Monte Carlo samples, our computed squared error values are random. We have investigated simultaneous perturbation stochastic

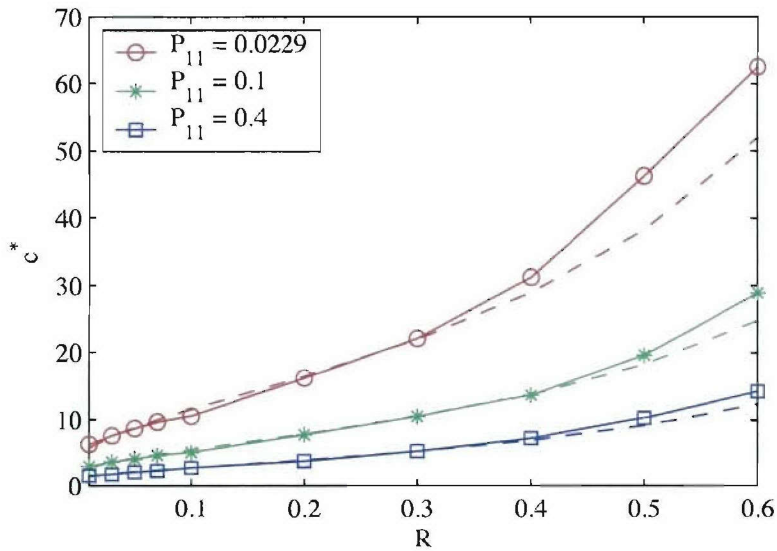

Figure 5: Heuristic vs. optimal values for c. The heuristics are plotted with solid lines; the optimal values are plotted with dashed lines. These values are plotted as a function of the observation noise variance R and the predicted position estimate error variance P_{11}.

approximation (SPSA), an iterative descent algorithm in which the gradient of the objective function is approximated by Monte Carlo samples, to find optimal values for the foveal center and foveal gain for myopic and non-myopic configuration.

Our SPSA results show that the myopic ($M = 1$) optimal value of c is a function of the observation noise variance R and the predicted position estimate error variance P_{11}. Using the SPSA computed optimal values of c as a model, we have developed a near-optimal heuristic function to compute c. Figure 5 compares the optimal values for c with the heuristic values.

Using Monte Carlo simulations, we evaluated the average squared error performance of SPSA in open-loop and open-loop feedback configurations. We also evaluated the performance of the heuristic and heuristic combined with SPSA in an approximate closed-loop configuration. Figure 6 shows the average squared error as a function of time for these different configuration approaches. The approximate closed-loop feedback approach and the myopic SPSA approach have the best performance. The multi-step ($M = 2$) open-loop and open-loop feedback approaches perform least well. The heuristic requires several orders of magnitude less computation than the SPSA methods.

4.2 Joint Tracking and Detection of Targets

In many scenarios, targets may be difficult to detect with a single use of a sensor due to target stealth, clutter, or other factors. The use of a tracker that incorporates a target dynamics model and multiple uses of the sensor may allow better target detection. Such as scheme, in which the estimator jointly tracks and make decisions about target presence or absence, is denoted track before detect. We have investigated joint tracking and detection in the context of the foveal sensor.

In previous work, a Bayesian approach to track before detect filtering for a single target was implemented using a discretization of the target state space. More recent work has has implemented this Bayesian approach using a particle filter.

Implementation of the track before detect filter for the foveal sensor required extension of existing work in two areas.

- To perform track before detect, we require a sensor that provides information without the necessity of target detection. The sensor model that we adopted is a foveal sensor that provides a one-dimensional image in which pixel values are related to the probability of the target being present at a given location. If the target is positioned such that its image falls in a given pixel, the pixel value's mean is related to the target intensity. The variance of

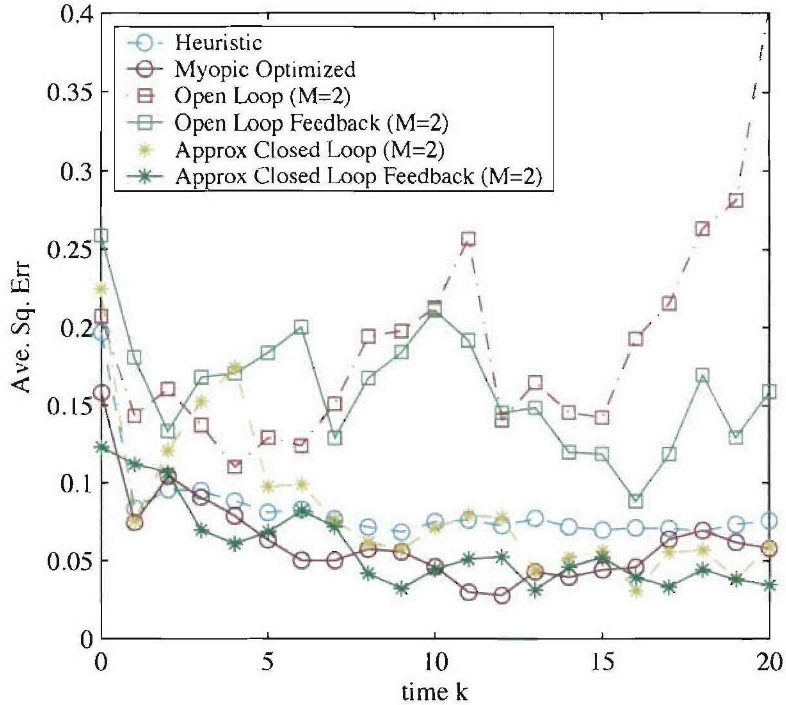

Figure 6: Average squared error as a function of time for the investigated configuration approaches.

all pixel values is determined by the variance of the imager noise. Such a model could result from an IR sensor, for example.

- To perform simultaneous tracking and detection, we must extend the state estimated by the particle filter to include both the number of targets (which may be zero) and the kinematic information for each target.

The use of the particle filter to implement the Bayesian track before detect scheme, while simple in concept, required significant work to achieve computationally efficient implementations. In particular, we devoted a significant amount of effort to the development of particle filters for pixelated sensors and for tracking an unknown and changing number of targets.

We adapted the near optimal sensor configuration heuristic to the pixelated foveal sensor model, and found that it did a good job of tracking targets that had a reasonably high signal-to-noise ratio. We were less successful simultaneously tracking and detecting dim targets; we believe that further development of the particle filter implementation may improve the performance.

4.3 Attentive Tracking of Closely Spaced Targets

We investigated heuristic strategies to configure a foveal sensor to track multiple targets moving in one dimension. These strategies were obtained by extending the heuristic strategy to multiple targets using three approaches: simultaneously observe all targets (SO), center the foveal region on each target in turn (TO), and center the foveal region on the target with the worst position estimate (WO). Our simulation results showed the best performance is obtained by the WO rule.

The dynamics model for each target and the model for the foveal sensor were the same as the models used for Task 2.1. The tracker was implemented using a joint multi-target probability density particle filter with joint probabilistic data association (JPDA). As part of this work, we developed novel independent partition (IP) proposal schemes to address the problem of association of observations with targets. In particular, we investigated two different methods: generalized nearest neighbor (IP-GNN) and joint probabilistic data association (IP-JPDA). IP-JPDA provided better performance in the presence of clutter.

Table 1: Percentage of runs that converged with no clutter

Proposal method	SO	TO	WO
IP-JPDA	90.5%	99.5%	99.7%
KP	88.5%	97.1%	96.4%

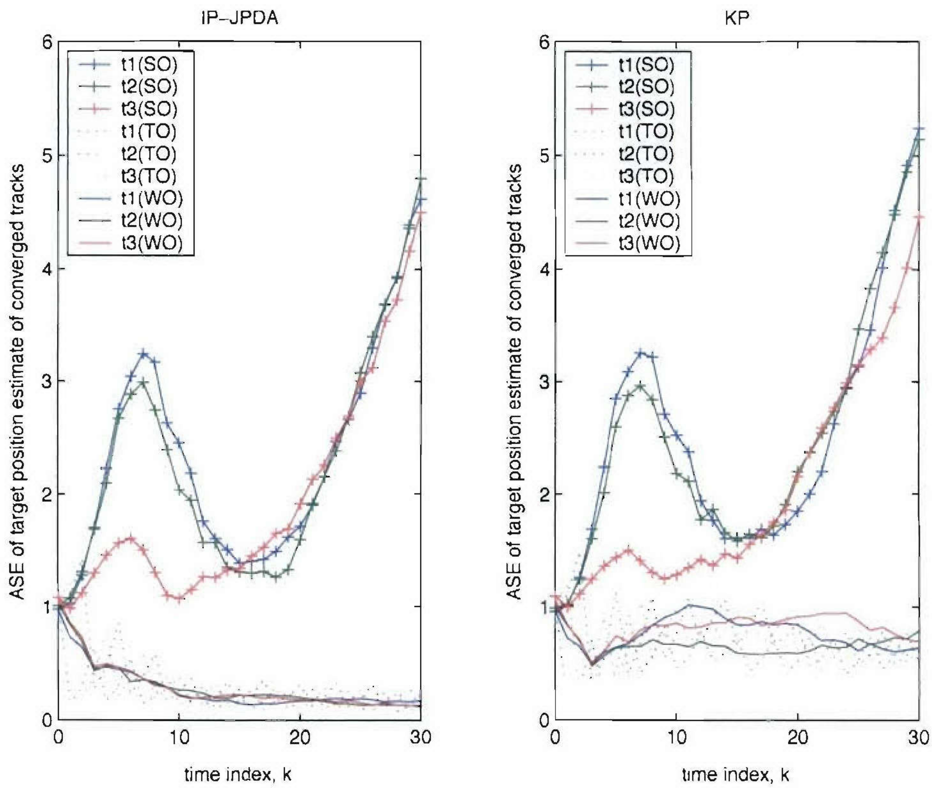

Figure 7: Average squared position error (ASE) for three targets (t1, t2, and t3) for the SO, TO and WO heuristics and no clutter.

We evaluated the foveal sensor performance in two scenarios:

- No clutter: the sensor detects all targets with probability one and there are no false alarms,
- Clutter: the sensor detects each target with probability 0.9 and false alarms are Poisson distributed with an expected number of false alarms of 0.5.

We compared the performance of the three configuration approaches using Monte Carlo simulation; all evaluations used 1000 simulation runs. We also compared the performance given by the IP-JPDA proposal distribution with the performance given by the kinematic prior (KP) proposal.[1] We characterized the performance by examining the

[1]The kinematic prior is the simplest proposal distribution to implement, and is often the first proposal distribution considered in a particle filter implementation.

Table 2: Percentage of runs that converged in clutter

Proposal method	SO	TO	WO
IP-JPDA	81.1%	90.7%	92.5%
KP	80.6%	75.8%	83.1%

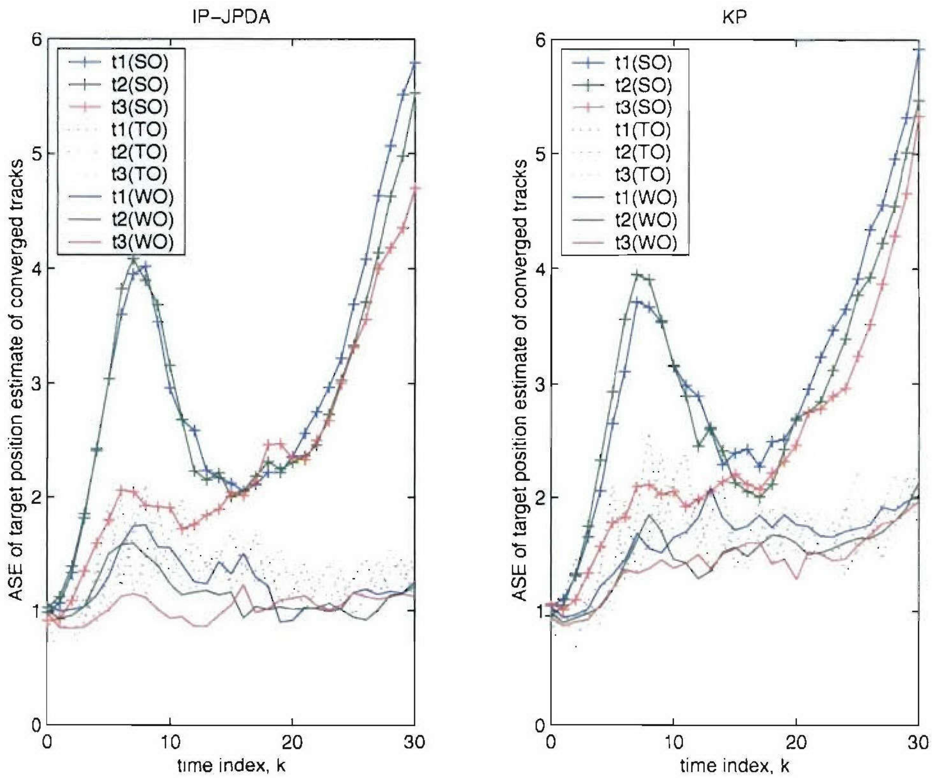

Figure 8: Average squared error (ASE) for the SO, TO and WO heuristics in clutter.

Table 3: Percentage of runs that converged using IP-JPDA and IP-GNN

Proposal method	no clutter	clutter
IP-JPDA	100%	98.9%
IP-GNN	100%	84.1%

percentage of Monte Carlo runs in which the tracker converged and, for converged runs, the average squared error in the estimated position of each target. We defined a run to have converged if the position errors for all targets remained below a given threshold during the last ten time steps of each run.

We compared the performance of the SO, TO, and WO approaches for both IP-JPDA and kinematic prior (KP) proposals for no clutter and clutter. Table 1 and Figure 7 show the percent of runs that converged and the average squared error (ASE) in the target position estimate as a function of time for the SO, TO and WO heuristics tracking three targets in no clutter. Table 2 and Figure 8 show the corresponding performance with clutter. The performance of IP-JPDA is always better than that of KP, and WO and TO with IP-JPDA proposal perform better than SO. Also, WO appears to be generally slightly better than TO in the presence of clutter.

We also compared IP-JPDA and IP-GNN. Table 3 and Figure 9 show that IP-JPDA has the same performance tracking 2 targets as IP-GNN in no clutter, while IP-JPDA has better performance than IP-GNN in clutter. This is because IP-JPDA is more robust with respect to false measurement-to-target assignments.

4.4 Configuration of Pan-ilt-zoom Cameras

We adapted the concepts develop for the foveal sensor configuration algorithms to the problem of configuring pan-tilt-zoom cameras to track a target maneuvering in three dimensions. In particular, we developed an adaptive zoom algorithm that minimizes localization errors by adaptively changing the focal length of the camera. Increasing the focal length (zooming in on a target) enhances the localization accuracy. However, increasing the focal length to much can result in a target loss when the target image falls off the image plane. Decreasing the focal length reduces the

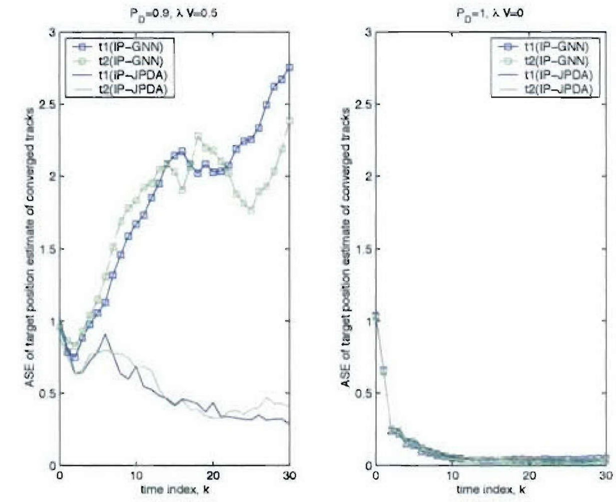

Figure 9: Comparison between IP-JPDA and IP-GNN proposal schemes for clutter and no clutter.

likelihood that the target will be lost but provides less accurate observations of target location.

The configuration algorithm uses a particle filter with a constant-velocity target dynamics model and a simple three-dimensional camera geometry model. The adaptive zoom algorithm adjusts the camera focal length until a given percentage of projected particles fall onto the image plane; the focal length is also adjusted by a confidence factor that influences and is influenced by whether the zooming algorithm is aggressive or conservative. The algorithm, which we called the Adaptive Zoom Technique for Enhanced Capture (AZTEC), uses two cameras to track a point target.

The algorithm was subsequently improved by using Monte Carlo simulation to establish the relationship between zoom and target tracking accuracy for different scenarios. This relationship was used to create a heuristic configuration algorithm with significantly improved tracking performance.

5 Task 3: Target Tracking Using Attentive Sensor Networks

This task was not addressed directly under AFOSR support. However, with funding from the DARPA ISP program, several different sensor configuration strategies for target tracking have been developed.

The first to be developed was for tracking a target as it moves through an array of very simple sensors; the sensors are modeled as energy detectors whose probability of detection varies inversely with the distance between the sensor and the target. The target is tracked using a particle filter. The sensors are configured to be either on or off using a simple heuristic: turn on the sensors that fall within a circle centered at the estimated target location and whose radius is determined by the uncertainty in the target position estimate. In addition to the simple heuristic, sensor configuration strategies based on the branch and bound techniques and integer programming have been developed.

In addition to sensor networks, the DARPA ISP program has funded work in which a moving target is tracked by a mobile sensor; at each time, the sensor can make one of a finite number of moves. We have investigated myopic (greedy) and non-myopic optimization approaches for sensor positioning using open-loop and open-loop feedback control strategies. Our results show that non-myopic optimization leads to better tracking performance.